The Journal of Operational Risk (15–38) Volume 6/Number 3, Fall 2011

A copula-based simulation model for supply portfolio risk

Halis Sak
Yeditepe University, Kayışdağı, 34755 Istanbul, Turkey;
email: halis.sak@gmail.com

Çağrı Haksöz
Sabancı School of Management, Sabancı University, Orhanlı, Tuzla 34956,
Istanbul, Turkey; email: cagrihaksoz@sabanciuniv.edu

A copula-based simulation model for supply portfolio risk in the presence of dependent breaches of contracts is introduced in this paper. We demonstrate our method for a supply-chain contract portfolio of commodity metals traded at the London Metal Exchange (LME). The analysis of spot price data on six LME commodity metals leads us to use a t-copula dependence structure with t-marginals and generalized hyperbolic marginals for the log returns. We also provide efficient simulation algorithms using importance sampling for the normal and t-copula dependence structure to quantify risk measures, supply-at-risk and conditional supply-at-risk. Numerical examples on a portfolio of six commodity metals demonstrate that our proposed method succeeds in decreasing the variance of the simulations. A numerical sensitivity analysis for the choice of the copula function is also provided. To the best of our knowledge, this is the first paper proposing efficient simulation algorithms on a supply-chain contract portfolio that has a copula-based dependence structure with generalized hyperbolic marginals.

1 INTRODUCTION

Commodity price risk is very important for firms that consume various commodity metals in their operations. A recent McKinsey CEO survey report by Gyorey *et al* (2010) notes that 37% of the CEO respondents stated that, over the next five-year period, they were unprepared for an increase in the volatility of commodity prices. Moreover, commodity price risk is exacerbated in the presence of breach-of-contract

The authors are grateful to the editor-in-chief, the associate editor, three anonymous referees, and Wolfgang Hörmann for their valuable suggestions that substantially improved the paper. They also acknowledge the Institute for Statistics and Mathematics at the Vienna University of Economics and Business for permitting the use of their computing cluster.

risk (Haksöz and Şimşek (2010)). Breach-of-contract risk is a fundamental operational risk classified under "Clients, Products, and Business Processes" as well as the "Execution, Delivery, and Process Management" categories of the Basel II framework (see, for example, Cruz (2002), Chernobai *et al* (2002), Haksöz and Kadam (2009) and Haksöz and Şimşek (2010) for further details on this type of operational risk). A breach of contract may occur for several reasons. It may be intentional, in that a supplier may prefer to take advantage of favorable spot market prices instead of selling via fixed-price contracts. Firms certainly do pay penalty charges if they breach contracts, which may somewhat compensate the financial loss for the other party. However, reputations become tarnished and strategic alliances get broken. There is a clear need to assess the potential severity of breach-of-contract risk.

In contrast to the level of interest in the practice, there is little research activity in this area of operational risk addressing breach-of-contract risk and the methods used to assess and hedge it. In a single buyer supplier model, Haksöz and Seshadri (2007) valued an American-type abandonment option, which models the breach of contract for a supplier, written in a long-term contract with a fixed penalty. To assess the portfolio risk of various commodity supply-chain contracts, Haksöz and Kadam (2009) provided a supply portfolio risk measurement tool based on the celebrated CreditRisk$^+$ model. Haksöz and Kadam (2009) coined the term supply-at-risk (SaR) and also presented risk-metric computations for a supplier portfolio of petrochemicals. In Haksöz and Kadam (2009), breach of contract was assumed to occur due to spot price evolution. However, Haksöz and Kadam (2009) did not address the dependency issue surrounding multiple breaches of contracts in the portfolio. Wagner *et al* (2009) presented a model for correlated supplier defaults (due to many financial-economic factors, not simply breaches of contract) with a copula dependence structure. Most recently, Haksöz and Şimşek (2010) provided a model for pricing bundled options (abandonment and price renegotiation option) in a supply-chain contract. This type of bundled option was shown to be valuable for mitigating breach-of-contract risk.

In this paper, building on Haksöz and Kadam (2009), we contribute to the literature by providing an efficient simulation method for supply portfolio risk assessment, where the supply-chain contracts in the portfolio have a dependence structure. The efficient simulation method is borrowed from Sak *et al* (2010) and is modified for our problem. Moreover, the algorithm given in Sak *et al* (2010) is designed for the t-copula dependence structure. We modify it so that it works for the normal copula as well. To demonstrate the value of our method, we study in detail a supply-contract portfolio with a copula dependence structure that is comprised of a number of commodity metals that are traded at the London Metal Exchange (LME). We further provide efficient algorithms in order to compute risk metrics such as SaR and conditional supply-at-risk (CSaR).

This paper is organized as follows. Section 2 presents the mathematical details of the supply portfolio model. Section 3 conducts the marginal distribution and copula fitting to commodity metal data. Sections 4 and 5 present the efficient simulation algorithms with importance sampling (IS) for tail loss probabilities and conditional expectations, which are used for calculating SaR and CSaR, respectively. We then present our numerical results and managerial insights on the commodity metal portfolio in Section 6. Finally, we conclude in Section 7.

2 THE MODEL

We assume that a buyer procures a variety of metal commodities from different global suppliers using long-term fixed-price contracts. These commodity metals are also traded at the LME. The market prices are known and there are liquid spot markets. During the contract duration, the suppliers can breach their supply contracts for any reason. An actual breach-of-contract event is exogenous and is not modeled in this paper.[1] Moreover, multiple dependent breaches of contracts can occur at the same time. We also assume that the buyer has to go to the spot market in case there is a breach of contract by the suppliers. That is, the buyer does not have alternative suppliers for the specific commodities purchased in this portfolio, apart from the spot market option. Even if there are potential backup suppliers, the price quoted for such emergency orders would closely follow the spot market price at that particular time. To that end, the buyer will be exposed to multiple spot market price risks under several dependent breaches of contracts. Hence, the buyer needs to assess the SaR and CSaR for such supply-contract portfolios in order to better manage its breach-of-contract risks.

Following the mathematical setting of Haksöz and Kadam (2009), we assume that the long-term contract price is equal to the median metal spot price in this supply portfolio without loss of generality. Furthermore, we also assume that there is a fixed penalty paid by the suppliers in case of breach of contract and that this penalty covers the transaction costs required to purchase commodities from the spot market for the buyer.[2] Basically, these assumptions help us to delineate the impact of spot price risk in a portfolio of contracts without considering the actual penalty and transaction cost data. Thus, we can write the individual risk exposure for the buyer at breach of

[1] In practice, one also may need to determine the explicit breach-of-contract probabilities. These probabilities may be affected by internal and external factors. Environmental disruptions such as earthquakes, floods and hurricanes may cause suppliers to breach their contracts. On the other hand, a firm may go bankrupt and be unable to fulfill the contract, and therefore have to breach it. This nontrivial problem is left for future research.

[2] As pointed out by a referee, buyers can also design floating penalty contracts that may mitigate the price risk. However, enforcing such a contract would be harder.

contract as follows:

$$\epsilon_i = \max\{0, Q_i(P_i - \bar{P}_i)\} \tag{2.1}$$

where Q_i is the contracted quantity of the metal i, P_i is the spot price for metal i, and $\bar{P}_i$ is the median spot price for metal i. We use the median spot price as a proxy for the long-term contract price of the metal procured.

For a number of contracts in the portfolio, $i = 1, \dots, n$, the total risk exposure for a number of potential breaches of contract can be expressed as follows:

$$R = \sum_{i=1}^{n} \epsilon_i \tag{2.2}$$

where ϵ_i is as given in (2.1).

In this expression, we only have the financial impact at breach, that is, the severity of the breach events. Note that this severity is driven only by the spot price risks. We assume that the log returns of n metals over a day follow an elliptical copula and that its dependence structure is described by the positive definite matrix Σ. L denotes the (lower-triangular) Cholesky factor of Σ satisfying $LL' = \Sigma$. We consider only the normal and t-copula alternatives for the elliptical copula function. We give the model, algorithms and numerical results for the normal and t-copula together in order to save space. When writing the model and algorithms, we only give the differences between the normal and the t-copula. The classical random return vector generation algorithm from the normal and t-copula starts with a vector Z of d independent and identically distributed (iid) standard normal variates that is then transformed into the correlated normal vector $\tilde{Z} = LZ$ (see Hörmann and Sak (2010) for a different generation algorithm from the t-copula). For the t-copula, we obtain the vector T from the multivariate t-distribution by generating a random variate Y from a chi-squared distribution with ν degrees of freedom (χ^2_ν) and calculating $T = \tilde{Z}/\sqrt{Y/\nu}$. The log-return vector $S = (S_1, S_2, \dots, S_n)'$ is then the result of the component-wise transform:

$$S_i = c_i G_i^{-1}(F(V_i)) \tag{2.3}$$

where F denotes the cumulative distribution function (CDF) of a standard normal distribution for the normal copula and the CDF of a t-distribution with ν degrees of freedom for the t-copula. We have $V_i = \tilde{Z}_i$ for the normal copula and $V_i = T_i$ for the t-copula. G_i denotes the CDF of the marginal distribution of the return of the ith metal and c_i denotes the volatility scaling parameter of the ith metal defined in (2.5). Then, given that P_{i0} is the spot price for metal i at time 0, the spot price of metal i at the end of a time horizon of m days is:

$$P_i = P_{i0} \prod_{j=1}^{m} \mathrm{e}^{S_i^{(j)}} \quad \text{with } S_i^{(j)} = c_i G_i^{-1}(F(V_i^{(j)})),\ j = 1, \dots, m \tag{2.4}$$

where c_i denotes a scaling factor related to the daily volatility σ_i and the variance var_i of the ith marginal distribution by the formula:

$$c_i = \sigma_i \sqrt{\frac{1}{\text{var}_i}} \tag{2.5}$$

First, we are interested in the SaR, which is the quantile of the total risk exposure given in (2.2) for a probability level. To compute the SaR for a time horizon we require an efficient simulation algorithm with which to compute $P(R > x)$ for different values of x (perhaps simultaneously). Then an inversion of the tail loss probability distribution can be used to compute the SaR. Second, we provide an algorithm to compute the conditional expectations, which can be used to compute the CSaR given that we have the SaR.

3 MARGINAL DISTRIBUTION AND COPULA FITTING TO COMMODITY METAL DATA

We use the inference functions for margins method to fit a dependence structure between log returns of LME daily metal cash price data[3] as it is a simple and efficient method (see Malevergne and Sornette (2006) and Karadağ (2008) for other parametric and nonparametric possibilities). In this method, as the first step, the parameters for marginal distributions are estimated using likelihood maximization, then the parameters of the copula are estimated, again using the maximum-likelihood method and the estimated marginal distributions in the first step.

Given that T daily log returns are available for metal i, the log-likelihood maximization problem for the first step is:

$$\max_{\beta_i} \sum_{t=1}^{T} \ln(f_i(x_i^t; \beta_i)), \quad i = 1, \dots, n \tag{3.1}$$

where f_i is the probability density function and β_i is the parameter vector for candidate distribution. As candidate distributions, we try three different alternatives: Gaussian, t-distribution with location and scale, and the generalized hyperbolic distribution. The number of parameters needing to be estimated for these continuous distributions is two, three and five, respectively. The t-distribution is a natural candidate for the marginal distribution as it is a simple extension of the Gaussian distribution. On the other hand, our motives for trying the generalized hyperbolic distribution come from the field of finance, in which the most flexible and best fitting distribution to financial

[3] London Metal Exchange historical metal price data for the current year is freely available at www.lme.com.

data seems to be generalized hyperbolic distribution (Aas and Haff (2006); Behr and Pötter (2009); and Prause (1997)).

In the second step, the log-likelihood function of the copula is:

$$\max_{\alpha} \sum_{t=1}^{T} \ln c(F_1(x_1^t; \hat{\beta}_1), \dots, F_n(x_n^t; \hat{\beta}_n); \alpha) \tag{3.2}$$

where F_i is the CDF of the marginal distribution for metal i, c is the density of copula function, and α is the parameter vector for the copula. We consider only the normal and t-copula alternatives since Kole *et al* (2007) conclude that, among other alternatives such as the normal and the Gumbel copula, the t-copula is the best fitting copula for the risk management of linear asset portfolios in finance.

We use R (R Development Core Team (2008)) as a convenient working environment for solving (3.1) and (3.2), and for carrying out our simulations in the next sections. We use the R packages "fitdistrplus" (Delignette-Muller *et al* (2010)) for fitting the Gaussian and the t-distribution, "ghyp" (Breymann and Lüthi (2008)) for fitting the generalized hyperbolic distribution and "copula" (Yan and Kojadinovic (2010)) for fitting copulas.

The best fitting criteria for marginal distributions and copulas is the magnitude of log-likelihood values. However, since it does not account for the estimated number of parameters, we also look at the Akaike information criterion (AIC) values. The higher the log-likelihood value and the smaller the AIC value, the better the fit is. The AIC value is calculated using "fitdistrplus" and the "ghyp" package, as in Matteis (2001):

$$\text{AIC} = -2 \times \text{LogLik} + 2 \times \text{NE} \tag{3.3}$$

where LogLik denotes the log-likelihood value and NE denotes the number of estimated parameters.

We use daily LME spot prices for copper (Cu), aluminum (Al), nickel (Ni), zinc (Zn), lead (Pb) and tin (Sn) for the year 2010, as shown in Figure 1 on the facing page. The first eleven months of the data were used to fit the dependence structure and the marginal distributions. The last month of data is used for measuring the goodness of fit of the marginal distributions. The correlation matrix of the log returns is given in Table 1 on page 22. The maximum linear correlation is 0.844, which is between copper and zinc. The minimum linear correlation is 0.496, which is between lead and tin.

We note that the number of metals and duration of the data is quite limited for deriving conclusions for all of the metal spot markets. In this section our aim is to see whether there is a tendency for the metal data to deviate from the normal distribution. If so, what is the best distribution fit? Also, does the t-copula dependence structure fit well?

FIGURE 1 Commodity metal spot prices in 2010.

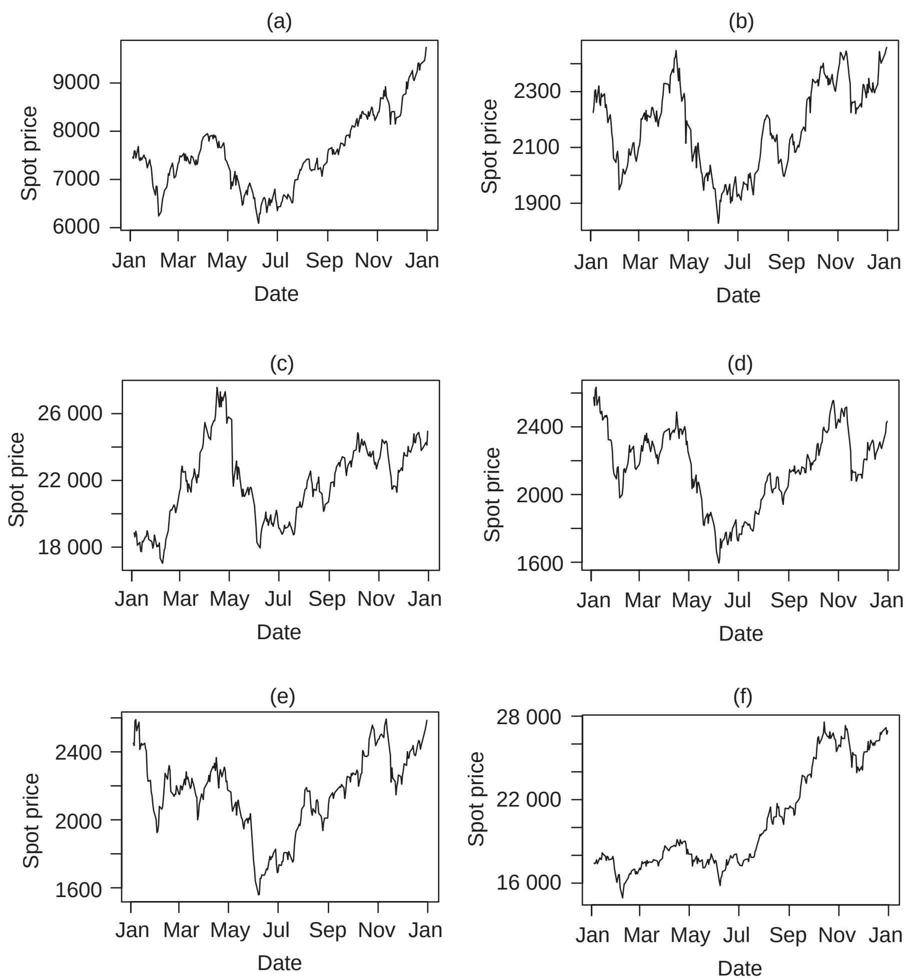

(a) Copper. (b) Aluminum. (c) Nickel. (d) Zinc. (e) Lead. (f) Tin.

Shapiro–Francia (SF), Anderson–Darling (AD), Cramer–von Mises (CVM), Lilliefors, and Pearson chi-squared tests are applied to test the normality of metal data using built-in functions in the "nortest" (Gross (2006)) R package (Ricci (2005)). The estimated p-values of the test statistics are given in Table 2 on the next page. Based on a significance level of 0.05, the null hypothesis stating that the data comes from the normal distribution is rejected in most of the tests. The SF, AD and CVM tests give more weight to the tails than the Lilliefors and the Pearson chi-squared tests. The

TABLE 1 Correlation matrix of the metal log returns.

	Cu	Al	Ni	Zn	Pb	Sn
Cu	1.000	0.785	0.697	0.844	0.595	0.678
Al	0.785	1.000	0.685	0.753	0.606	0.621
Ni	0.697	0.685	1.000	0.664	0.553	0.583
Zn	0.844	0.753	0.664	1.000	0.679	0.657
Pb	0.595	0.606	0.553	0.679	1.000	0.496
Sn	0.678	0.621	0.583	0.657	0.496	1.000

TABLE 2 p-values for five normality tests for metal log returns.

Metal	SF test	AD test	CVM test	Lilliefors test	Pearson test
Cu	0.043	0.026	0.028	0.093	0.241
Al	0.001	0.010	0.020	0.118	0.408
Ni	0.000	0.002	0.002	0.011	0.067
Zn	0.002	0.001	0.001	0.003	0.130
Pb	0.000	0.000	0.000	0.002	0.509
Sn	0.000	0.001	0.002	0.002	0.179

p-values suggest that log returns for copper could be assumed to be normal, although the SF, AD and CVM tests oppose this. However, the semi-heavy tails of log returns of the other metals indicate that the log returns do not follow a normal distribution.

Motivated by the nonnormality of the log returns, we fit the t-distribution and the generalized hyperbolic distribution. The log-likelihood and AIC values for the Gaussian, t- and generalized hyperbolic distribution fits are given in Table 3 on the facing page. In all of the cases, the generalized hyperbolic distribution produces the highest log-likelihood value. However, it is a five-parameter distribution. Thus, the AIC values are not always the minimum. It is better to use the generalized hyperbolic distribution for the first four metals (Cu, Al, Ni and Zn) and it is better to use t-distribution for the last two (Pb and Sn). The estimated parameters of the fitted marginal distributions for the log returns and p-values for the Kolmogorov–Smirnov (KS) test for measuring the goodness of fit are given in Table 4 on the facing page. We use log-return data for the month of December to compute the KS test statistics since this statistic cannot be used when parameters of the distributions need to be estimated from the data. The AD, CVM and Lilliefors tests are all modifications of the KS test. Computed p-values for the KS test are greater than 0.05, which leads to the conclusion that we cannot reject the null hypothesis that log returns in December come from those marginal

TABLE 3 Log-likelihood and AIC values for Gaussian, t- and generalized hyperbolic distributions for metals.

	Gaussian		t		Generalized hyperbolic		
Metal	**Log-likelihood**	**AIC**	**Log-likelihood**	**AIC**	**Log-likelihood**	**AIC**	**Which to use?**
Cu	606.93	−1209.86	607.14	−1208.28	610.24	−1210.49	GH
Al	615.82	−1227.65	619.04	−1232.08	622.57	−1235.15	GH
Ni	543.16	−1082.33	549.12	−1092.24	555.03	−1100.05	GH
Zn	544.95	−1085.90	545.79	−1085.59	551.08	−1092.15	GH
Pb	525.11	−1046.21	539.64	−1073.28	541.25	−1072.49	t
Sn	587.92	−1171.84	592.63	−1179.25	594.51	−1179.02	t

In the last column, 'GH' stands for generalized hyperbolic.

TABLE 4 Parameters of the fitted marginal distributions for metal log returns and p-values for the KS test.

	Generalized hyperbolic					t			
Metal	λ	α	δ	β	μ	**Location**	**Scale**	**df**	**KS**
Cu	7.1683	254.62	0.0002	−83.39	0.0212	—	—	—	0.051
Al	−2.8473	121.55	0.0399	−62.05	0.0146	—	—	—	0.365
Ni	0.6901	269.48	0.0421	−188.01	0.0479	—	—	—	0.630
Zn	1.5894	123.22	0.0259	−48.81	0.0192	—	—	—	0.532
Pb	—	—	—	—	—	0.00084	0.0187	4.69	0.148
Sn	—	—	—	—	—	0.00228	0.0154	5.46	0.294

distributions. We use alpha/delta parametrization (Breymann and Lüthi (2008)) for the generalized hyperbolic distribution to print the estimated parameters, as we use this parametrization in our simulation functions. For the t-distribution, there are three parameters estimated, ie, location, scale and degrees of freedom (df).

The histograms of the log returns with the fitted t- and generalized hyperbolic distributions, and Q–Q plots for only copper and aluminum are given in Figure A.1 on page 33 and Figure A.2 on page 34.[4] The log returns for copper seem to be quite close to the normal distribution. This visual observation is consistent with the tabulated p-values in Table 2 on the facing page. However, for the other five metals, the t- and generalized hyperbolic distributions capture the high kurtosis and fat tails of the data.

[4] Plots for the rest of the commodity metals are available from the authors upon request.

TABLE 5 Results of copula fitting for a portfolio consisting of all metals.

Copula	Parameter(s)	Standard error	Log-likelihood	AIC
Normal	ρ_{norm}	$\text{SE}_{\rho_{\text{norm}}}$	505.64	−981.28
Student t	ρ_t, $\upsilon = 11.53$	SE_{ρ_t}, $\text{SE}_{\upsilon} = 2.45$	523.86	−1015.72

TABLE 6 Correlation matrix of the fitted normal copula (ρ_{norm}) for metal log returns.

	Cu	Al	Ni	Zn	Pb	Sn
Cu	1.000	0.774 (0.021)	0.683 (0.029)	0.834 (0.015)	0.619 (0.034)	0.670 (0.030)
Al	0.774	1.000	0.667 (0.030)	0.741 (0.023)	0.620 (0.034)	0.611 (0.034)
Ni	0.683	0.667	1.000	0.641 (0.032)	0.569 (0.038)	0.573 (0.037)
Zn	0.834	0.741	0.641	1.000	0.703 (0.027)	0.657 (0.031)
Pb	0.619	0.620	0.569	0.703	1.000	0.524 (0.041)
Sn	0.670	0.611	0.573	0.657	0.524	1.000

Standard errors ($\text{SE}_{\rho_{\text{norm}}}$) are given in parentheses.

We fit elliptical copulas, the normal and t-copula, to the log returns of the metal data using the estimated marginals. As the dimension of the portfolio increases, the expression of the probability density functions for the Archimedean copulas become more complex and thus the probability density function is not available due to the intensive computing involved in differentiating the CDF (see Yan (2007) and Karadağ (2008)). Moreover, the t-copula is preferred to the Gaussian and Gumbel copulas because of its ability to capture the dependence better in the nonextreme and extreme (tails) of financial returns.

For copula fitting, we use built-in functions in the "copula" (Yan and Kojadinovic (2010)) R package. Fitting results are summarized in Table 5. The correlation matrix and standard error of the point estimates are given for normal and t-copulas in Table 6 and Table 7 on the facing page, respectively. The numerical results given in Table 5 suggest that the t-copula is better than the normal copula at capturing the dependence structure of metal log returns.

To conclude this section, according to the empirical results from the limited data, the t-copula with t-marginals and generalized hyperbolic marginals seems to be an

TABLE 7 Correlation matrix of the fitted t-copula (ρ_t) for metal log returns.

	Cu	Al	Ni	Zn	Pb	Sn
Cu	1.000	0.779 (0.023)	0.700 (0.030)	0.833 (0.017)	0.675 (0.033)	0.659 (0.034)
Al	0.779	1.000	0.672 (0.032)	0.741 (0.026)	0.662 (0.033)	0.602 (0.039)
Ni	0.700	0.672	1.000	0.653 (0.034)	0.599 (0.039)	0.574 (0.041)
Zn	0.833	0.741	0.653	1.000	0.747 (0.026)	0.647 (0.035)
Pb	0.747	0.662	0.599	0.747	1.000	0.559 (0.042)
Sn	0.659	0.602	0.574	0.647	0.559	1.000

Standard errors (SE_{ρ_t}) are given in parentheses.

adequate model with which to capture the dependencies and to explain the semi-heavy tails of the returns of commodity metal data.

4 SIMULATING TAIL LOSS PROBABILITIES USING IMPORTANCE SAMPLING

Algorithm A.1 gives all the details of the naive simulation algorithm necessary to evaluate the tail loss probability $P(R > x)$ for a time horizon of m days.

We modify the IS technique described in Sak *et al* (2010) for our problem. The technique is summarized as follows. We add a mean shift vector with positive entries to the normal vector Z for the normal copula and additionally use a scale parameter $\theta < 2$ for the gamma (scale parameter two corresponds to chi-squared distribution) random variate Y in order to increase the probability of very high returns for the t-copula. The main practical problem in the application of IS is the choice of the parameters of the IS distribution. We only give the algorithms here. For a better understanding of the technique we refer the interested reader to Sak *et al* (2010).

We use the R package "Runuran" (Leydold and Hörmann (2008)) for evaluating quantiles from the generalized hyperbolic distribution (see Sak *et al* (2010, Section 3)). "Runuran" uses a numerical inversion algorithm that requires only the probability density function instead of the cumulative density function (Derflinger *et al* (2009, 2010)).

Algorithm A.3 returns the optimal mean-shift μ for the normal copula and also the optimal mode y_0 of the IS density for Y for the t-copula. Note that function $R()$ used in step 2 of Algorithm A.2 denotes the total risk exposure defined in (2.2). The following equation gives the optimal scale parameter θ of the gamma IS density for Y:

$$\theta = \frac{y_0}{\frac{1}{2}\nu - 1}. \tag{4.1}$$

The likelihood ratio for the normal copula is:

$$W_\mu(Z) = \exp(-\mu' Z + \tfrac{1}{2}\mu'\mu), \tag{4.2}$$

and for the t-copula it is:

$$W_{\mu,\theta}(Z, Y) = \exp(-\mu' Z + \tfrac{1}{2}\mu'\mu - \tfrac{1}{2}Y + Y/\theta + \log(\tfrac{1}{2}\theta)\tfrac{1}{2}\nu) \tag{4.3}$$

where $\exp(-\mu' Z + \frac{1}{2}\mu'\mu)$ accounts for the mean shift that we have added to the normal vector, and the term $\exp(-\frac{1}{2}Y + Y/\theta + \log(\frac{1}{2}\theta)\frac{1}{2}\nu)$ relates the density of the chi-squared distribution with degrees of freedom ν to that of the gamma distribution with shape parameter $\frac{1}{2}\nu$ and scale parameter θ. The final IS algorithm is presented in Algorithm A.4.

The SaR associated with probability $1 - \alpha$ is the quantile:

$$\text{SaR}_\alpha = \inf\{x \colon P(R > x) \leqslant \alpha\} \tag{4.4}$$

Algorithm A.4 can be used to simulate tail loss probabilities for various threshold levels to derive a probability distribution of the total risk exposure. Then a regression algorithm can be used to calculate SaR_α.

5 SIMULATING CONDITIONAL EXPECTATIONS USING IMPORTANCE SAMPLING

In this section we tackle the problem of simulating conditional expectation $E[R \mid R > x]$.

If we assume that $P(R > x) > 0$, $E[R \mid R > x]$ can be written as:

$$r = E[R \mid R > x] = \frac{E[R\mathbf{1}_{\{R>x\}}]}{P(R > x)} \tag{5.1}$$

where $\mathbf{1}_{\{\cdot\}}$ denotes the indicator function.

The naive simulation estimate for this ratio is:

$$\hat{r}^{\text{naive}} = \frac{\sum_{k=1}^{N} R^{(k)}\mathbf{1}_{\{R^{(k)}>x\}}}{\sum_{k=1}^{N} \mathbf{1}_{\{R^{(k)}>x\}}} \tag{5.2}$$

To estimate the accuracy of (5.2), we use the $\delta\%$ confidence interval (Glasserman (2004, 2005)):

$$\hat{r}^{\text{naive}} \pm z_{\delta/2}\frac{\hat{\sigma}^{\text{naive}}}{\sqrt{N}} \tag{5.3}$$

where:

$$\hat{\sigma}^{\text{naive}} = \left(\frac{N \sum_{k=1}^{N} (R^{(k)} - \hat{r}^{\text{naive}})^2 \mathbf{1}_{\{R^{(k)} > x\}}}{(\sum_{k=1}^{N} \mathbf{1}_{\{R^{(k)} > x\}})^2} \right)^{1/2} \tag{5.4}$$

and $z_{\delta/2}$ denotes the quantile of the standard normal distribution for the probability level of $\delta/2$.

Algorithm A.1 gives details on how to use the naive simulation estimate.

Following Glasserman (2005) and Sak and Hörmann (2011), we use the IS distribution computed for the problem $(P(R > x))$ in simulating $E[R \mid R > x]$. The IS simulation estimate and its $\delta\%$ confidence interval (see Glasserman (2004, 2005)) is as follows:

$$\hat{r}^{\text{IS}} = \frac{\sum_{k=1}^{N} R^{(k)} W^{(k)} \mathbf{1}_{\{R^{(k)} > x\}}}{\sum_{k=1}^{N} W^{(k)} \mathbf{1}_{\{R^{(k)} > x\}}} \tag{5.5}$$

and:

$$\hat{r}^{\text{IS}} \pm z_{\delta/2} \frac{\hat{\sigma}^{\text{IS}}}{\sqrt{N}} \tag{5.6}$$

where:

$$\hat{\sigma}^{\text{IS}} = \left(\frac{N \sum_{k=1}^{N} (R^{(k)} W^{(k)} - \hat{r}^{\text{IS}} W^{(k)})^2 \mathbf{1}_{\{R^{(k)} > x\}}}{(\sum_{k=1}^{N} W^{(k)} \mathbf{1}_{\{R^{(k)} > x\}})^2} \right)^{1/2} \tag{5.7}$$

and $z_{\delta/2}$ denotes the quantile of the standard normal distribution for the probability level of $\delta/2$.

Details of how to use this estimate are presented as Algorithm A.4. Using this algorithm we can compute $\text{CSaR}_\alpha = E[R \mid R > \text{SaR}_\alpha]$.

6 NUMERICAL RESULTS

We use the IS algorithms given in Sections 4 and 5 for simulating the total risk exposure of the metal portfolio analyzed in Section 3. The log returns following the t-copula with generalized hyperbolic or t-marginals is the dependence structure the data suggests. We use the fitting results summarized in Table 5 on page 24 and Table 7 on page 25 for the t-copula. The correlation matrices (Σ that we use in the simulation algorithms) for the normal and t-copula are given in Table 6 on page 24 and Table 7 on page 25. For the numerical results presented we use $Q_i = 1, i = 1, \dots, n$.

The efficiency of a simulation method is inversely proportional to the product of the sampling variance and the required simulation time. We therefore report, as a main result of our comparison efficiency ratio, the ratio of the product of the sampling variance and the execution time of the naive method and the IS method.

TABLE 8 SaR_α, CSaR_α values and CSaR_α 95% confidence interval as percentage of the point estimates for the naive and the IS methods over a one-day horizon.

α	SaR_α	CSaR_α	95% confidence interval: naive method	95% confidence interval: IS method	Efficiency ratio
0.05	9839.5	10341.6	±0.14%	±0.04%	10
0.01	10635.7	11182.3	±0.34%	±0.05%	46
0.005	10989.1	11573.8	±0.48%	±0.05%	75
0.002	11489.6	12147.0	±0.74%	±0.06%	148
0.001	11902.7	12638.7	±1.31%	±0.06%	382

We use the IS algorithm given in Algorithm A.4 to compute tail loss probabilities for various threshold levels. We then fit a cubic smoothing spline to tail loss probabilities versus thresholds data in order to compute SaR_α for a number of α values in Table 8. We use the R package "fields" (Furrer *et al* (2010)) for fitting a cubic smoothing spline to the data. We also provide CSaR_α values and the half-length of the 95% confidence intervals in percentages for CSaR_α. The efficiency ratios indicate the relative efficiency of the IS with respect to the naive simulation in computing CSaR_α values. The efficiency ratios increase as the event simulated becomes rarer. This is an attribute of IS. Execution times are 17.0 and 19.7 seconds, respectively, for the naive and the IS methods for computing CSaR_α values for $N = 100\,000$. Furthermore, the efficiency of the IS in computing tail loss probabilities with respect to the naive simulation is quite similar to the presented efficiency ratios for CSaR_α.

It is important to assess the sensitivity of the numerical results given in Table 8 to the choice of the copula function while keeping the marginal distributions and the correlation matrix the same as suggested in Johnson and Tenenbein (1981) for a similar problem. We use different degrees of freedom for the t-copula and the normal copula to see how the choice of the copula-based joint distribution affects the simulated results for the SaR_α and CSaR_α measures. Over a one-day horizon, SaR_α and CSaR_α values for sets of α and degrees of freedom of the t-copula ($\nu = \infty$ is the normal copula case) are provided in Table 9 on page 30. In particular, simulated results change very little; the maximum difference in these results is 3.2% for the ten cases considered. The fact that these results change very little adds credibility to the measures that we have developed.

As the degrees of freedom for the t-copula increase (approaching the normal copula), SaR_α and CSaR_α decrease for the tails ($\alpha < 0.05$). This is an expected result since the tail dependence between contracts is lessening (probability distribution of total risk is less fat). However, as we approach the center of the distribution, SaR_α and CSaR_α values increase as the dependence structure gets stronger. This is observed for $\mathrm{SaR}_{\alpha=0.05}$.

To give a rough idea of how the tail loss probabilities and CSaR_α change in time, we draw tail loss probabilities and CSaR_α of total exposures for time horizons of one day and one week simultaneously in Figure 2 on page 31 and Figure 3 on page 31. We use the IS and naive simulation for the one-day horizon and only the IS for the one-week horizon in computing the tail loss probabilities, CSaR_α and their confidence intervals. The efficiency of the IS method for the one-day horizon can be easily observed when we compare it with the naive simulation. Although we use a greater number of replications for the naive simulation, it gives wider confidence intervals and stops giving sensible confidence intervals for thresholds greater than 12 500.

The three curves show the sample mean and its 95% confidence interval.

When we compare the tail loss probabilities for one-day and one-week horizons, we observe that the tail loss probabilities increase as the time horizon is increased. For a time horizon of one day we use $m = 1$ in Algorithm A.4. For a one-week horizon we use $m = 5$ (five working days is equivalent to one week). Due to the complicated return function in step 2 of Algorithm A.2, as we extend the time horizon, the efficiency of the IS method decreases, as can be observed in Figure 2 on page 31 and Figure 3 on page 31. Although the wideness of confidence intervals for time horizons of one day and one week seem to be nearly the same in Figure 2 on page 31, the tail loss probability axis is given in a logarithmic scale in Figure 2 on page 31. Indeed, the correct way of simulating the tail loss probabilities and conditional expectations for long horizons such as the one-week horizon is to fit the marginal distributions and copulas for weekly instead of daily log returns. Then, after adjusting the scaling factor (σ_i should be this time weekly volatility), the presented algorithms could be used to simulate the tail loss probabilities and conditional expectations of total risk exposures for the one-week time horizon using $m = 1$.

For the same α level, the one-week horizon CSaR_α is greater than for the one-day horizon CSaR_α, in line with our observations in Figure 2 on page 31. Furthermore, confidence intervals get worse in computing one-day and one-week horizons CSaR_α as α decreases, as can be observed in Figure 3 on page 31. As α decreases, the likelihood ratios for the IS method decrease to make the losses equal to the threshold on average. This decrease in the likelihood ratios is responsible for the degradation in quality of the confidence intervals for the IS method.

7 CONCLUSION

In this paper we have introduced an efficient simulation model for quantifying the risk measures for supply portfolio risk in the presence of dependent breaches of contracts.

TABLE 9 SaR_α and CSaR_α values for sets of α and degrees of freedom of the t-copula over a one-day horizon ($\nu = \infty$ is the normal copula case).

ν	$\mathbf{SaR}_\alpha$	$\mathbf{CSaR}_\alpha$
$\alpha = 0.05$		
3	9 826.7	10 374.5
5	9 833.4	10 362.8
10	9 837.7	10 349.6
15	9 839.6	10 345.2
∞	9 840.0	10 326.1
$\alpha = 0.01$		
3	10 688.9	11 287.2
5	10 666.7	11 247.4
10	10 640.1	11 194.8
15	10 630.1	11 167.1
∞	10 602.1	11 109.6
$\alpha = 0.005$		
3	11 078.3	11 715.9
5	11 039.4	11 660.8
10	10 998.2	11 588.1
15	10 979.6	11 559.2
∞	10 933.5	11 475.7
$\alpha = 0.002$		
3	11 619.2	12 322.2
5	11 567.3	12 260.6
10	11 499.1	12 160.4
15	11 470.5	12 118.3
∞	11 397.0	12 007.9
$\alpha = 0.001$		
3	12 065.3	12 845.7
5	12 000.5	12 766.7
10	11 914.4	12 662.8
15	11 876.2	12 591.3
∞	11 774.6	12 448.6

FIGURE 2 Tail loss probabilities of total risk exposure for time horizons of one day and one week using the IS method (with 1000 replications) and the naive method (with 10 000 replications).

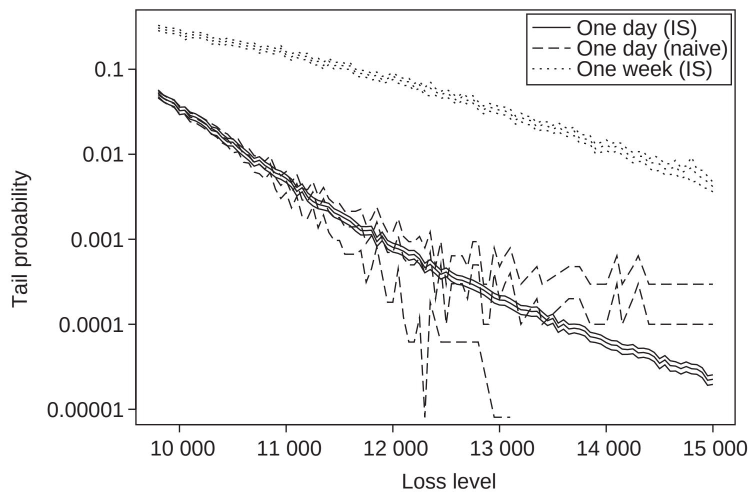

FIGURE 3 CSaR_α for time horizons of one day and one week using the IS method (with 1000 replications) and the naive method (using 10 000 replications).

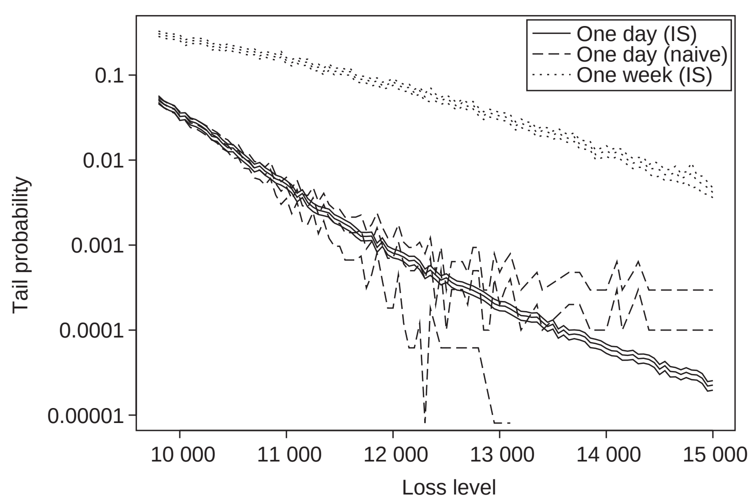

The three curves show the sample mean and its 95% confidence interval.

The model is based on a copula dependence structure. For assessing model parameters, we analyzed a limited data set of LME commodity metal spot prices. This process revealed a better fit for the t-copula dependence structure with t- or generalized hyperbolic marginal distributions for the log returns of the metals. Furthermore, we adopted the IS strategy given in Sak *et al* (2010) to compute SaR and CSaR under the normal and t-copula dependence structure. Our numerical results showed that the proposed method is much more efficient than a naive simulation for computing tail loss probabilities and conditional expectations. We also provided a numerical sensitivity analysis for the choice of the copula function.

The method proposed in this paper could very well assist supply-chain, procurement, and operational risk executives in assessing supply portfolio risk with a dependence structure. Future research is certainly required in developing models with explicit breach-of-contract probabilities and more sophisticated dependence structures. Furthermore, using real-life penalty and transaction cost data would be useful for quantifying the aggregate supply risk. We hope that our paper will motivate more research in this growing field.

APPENDIX A

A.1 Algorithm: computation of $P(R > x)$ and $E[R \mid R > x]$ using naive simulation for the normal and the t-copula

(0) Initialization:

- (a) Compute Cholesky factor L of Σ, ie, $LL' = \Sigma$.
- (b) Compute c_i for $i = 1, \dots, n$ using (2.5).

(1) Repeat for replications $k = 1, \dots, N$:

- (a) Repeat for replications $j = 1, \dots, m$:
  - (i) Generate independent standard normal variates Z then compute $\tilde{Z} = LZ$.
  - (ii) Generate Y from χ^2_ν-distribution for the t-copula.
  - (iii) $V_i = \tilde{Z}_i$ for the normal copula and $V_i = \tilde{Z}_i/\sqrt{Y/\nu}$ for the t-copula for $i = 1, \dots, n$.
  - (iv) Calculate $S_i^{(j)}$ for $i = 1, \dots, n$ using (2.3).
- (b) Calculate P_i for $i = 1, \dots, n$ using (2.4) then total risk exposure $R^{(k)}$ using (2.2).

(2) Return $(1/N)\sum_{k=1}^{N} \mathbf{1}\{R^{(k)} > x\}$ for computing $P(R > x)$ and return $\hat{r}^{\text{naive}}$ using (5.2) for computing $E[R \mid R > x]$.

FIGURE A.1 (a) Histogram and (b) generalized hyperbolic Q–Q plot for the log returns of copper.

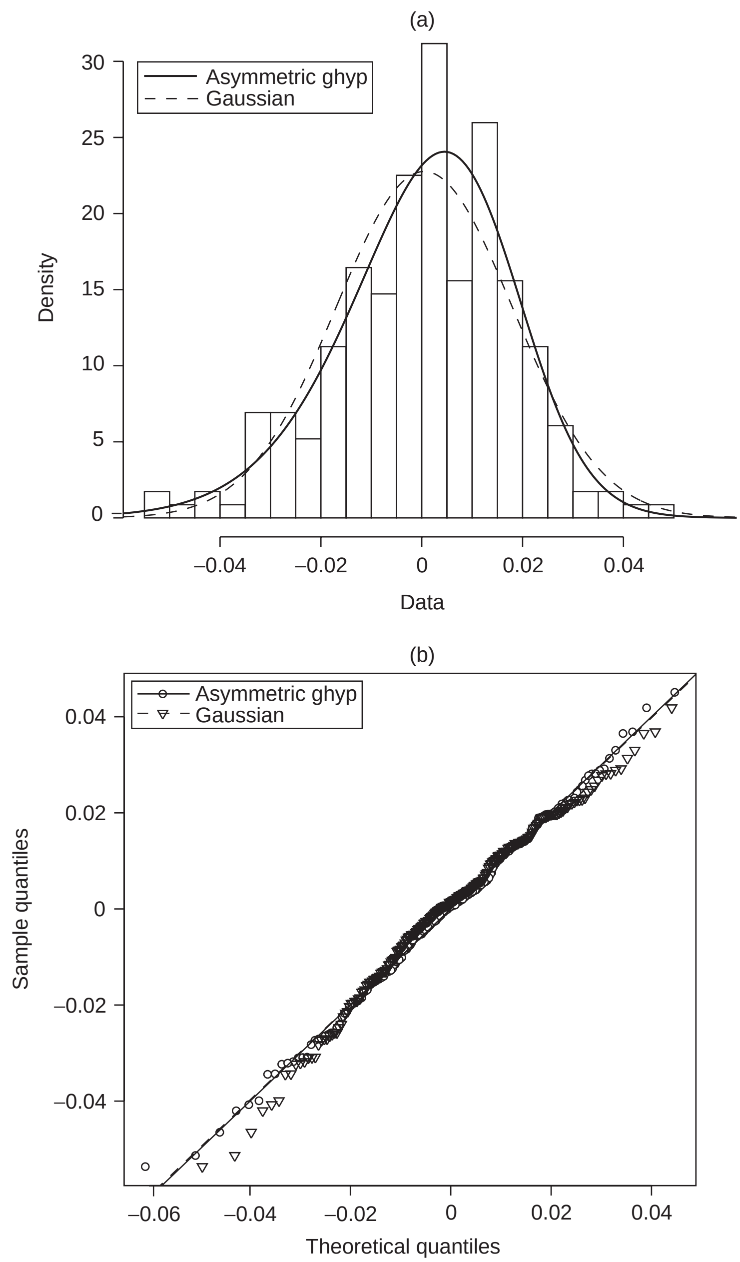

FIGURE A.2 (a) Histogram and (b) generalized hyperbolic Q–Q plot for the log returns of aluminum.

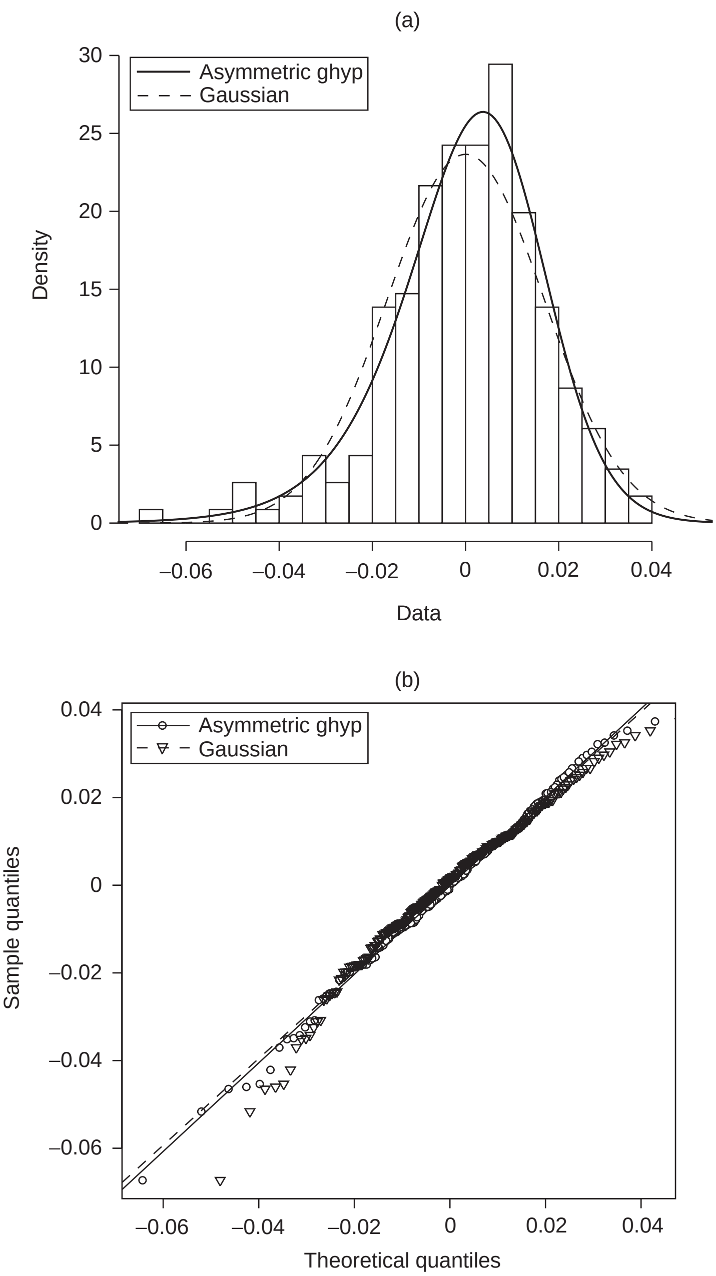

A.2 Algorithm: computation of z_0, y_0 and o_f for a given direction z_d

(1) Set $z_d^1 = z_d/\|z_d\|$.

(2) Compute r_0 by solving $(R(z = r_0 z_d^1) - (x + \Delta) = 0)$ numerically (use, for example, $\Delta = 10^{-5}$) for the normal copula and by solving $(R(z = r_0 z_d^1, y = \nu) - (x + \Delta) = 0)$ for the t-copula.

(3) Return vector $z_0 = r_0\, z_d^1$, and objective function value $o_f = -r_0^2$ for the normal copula and vector $z_0 = r_0 \sqrt{y_0/\nu}\, z_d^1$, $y_0 = (\nu - 2)/(1 + r_0^2/\nu)$, and objective function value $o_f = (\nu/2 - 1)(\log(y_0) - 1)$ for the t-copula.

A.3 Algorithm: computation of the mean shift vector μ and y_0

(0) Initialization:

(a) Compute Cholesky factor L of Σ, ie, $LL' = \Sigma$.

(b) Compute c_i for $i = 1, \dots, n$ using (2.5).

(1) Compute $z_d = L'c$.

(2) Call an optimization algorithm with starting direction z_d, objective function as given in Algorithm A.2, and nonnegativity constraints for all components of z_d (we used a quasi-Newton method with constraints). Get optimal direction ${}_{\text{opt}}z_d$.

(3) Call Algorithm A.2 with direction ${}_{\text{opt}}z_d$ and get the optimal vector z_0 for the normal copula, z_0 and y_0 for the t-copula.

(4) Return the optimal mean shift $\mu = z_0$ for the normal copula and $\mu = z_0$ and optimal mode y_0 for Y for the t-copula.

A.4 Algorithm: computation of $P(R > x)$ and $E[R \mid R > x]$ using IS for the normal and t-copula

(0) Initialization:

(a) Compute Cholesky factor L of Σ, ie, $L\,L' = \Sigma$.

(b) Compute c_i for $i = 1, \dots, n$ using (2.5).

(c) Compute μ for the normal copula and μ and y_0 for the t-copula using Algorithm A.3.

(d) Compute $\theta = y_0/(\nu/2 - 1)$ for the t-copula.

(1) Repeat for replications $k = 1, \dots, N$:

(a) Repeat for replications $j = 1, \dots, m$:

(i) Generate $Z_i \sim N(\mu_i, 1), i = 1, \dots, n$, independently, then compute $\tilde{Z} = LZ$.

(ii) Generate Y from gamma distribution with shape parameter $\frac{1}{2}\nu$ and scale parameter θ for the t-copula.

(iii) Calculate $W_{\mu}^{(j)}$ for the normal copula as in (4.2) and $W_{\mu,\theta}^{(j)}$ as in (4.3) for the t-copula.

(iv) $V_i = \tilde{Z}_i$ for the normal copula and $V_i = \tilde{Z}_i/\sqrt{Y/\nu}$ for the t-copula for $i = 1, \dots, n$.

(v) Calculate $S_i^{(j)}$ for $i = 1, \dots, n$ using (2.3).

(b) Calculate P_i for $i = 1, \dots, n$ using (2.4) then total risk exposure $R^{(k)}$ using (2.2).

(c) Calculate:

$$W^{(k)} = \prod_{j=1}^{m} W_{\mu}^{(j)}$$

for the normal copula and:

$$W^{(k)} = \prod_{j=1}^{m} W_{\mu,\theta}^{(j)}$$

for the t-copula.

(2) Return $(1/N)\sum_{k=1}^{N} W^{(k)}\mathbf{1}\{R^{(k)} > x\}$ for computing $P(R > x)$ and return $\hat{r}^{\mathrm{IS}}$ using (5.5) to compute $E[R \mid R > x]$.

REFERENCES

Aas, K., and Haff, I. H. (2006). The generalized hyperbolic skew Student's t-distribution. *Journal of Financial Econometrics* **4**(2), 275–309.

Behr, A., and Pötter, U. (2009). Alternatives to the normal model of stock returns: Gaussian mixture, generalised logf and generalised hyperbolic models. *Annals of Finance* **5**(1), 49–68.

Breymann, W., and Lüthi, D. (2008). ghyp: a package on generalized hyperbolic distributions. R package, version 1.5.4.

Chernobai, A. S., Rachev, S. T., and Fabozzi, F. J. (2002). *Operational Risk: A Guide to Basel II Capital Requirements, Models, and Analysis*. John Wiley & Sons.

Cruz, M. G. (2002). *Modeling, Measuring and Hedging Operational Risk*. John Wiley & Sons.

Delignette-Muller, M.-L., Regis Pouillot, J.-B. D., and Dutang, C. (2010). fitdistrplus: help to fit of a parametric distribution to non-censored or censored data. R package, version 0.1-4.

Derflinger, G., Hörmann, W., Leydold, J., and Sak, H. (2009). Efficient numerical inversion for financial simulations. In *Monte Carlo and Quasi-Monte Carlo Methods 2008*, L'Ecuyer, P., and Owen, A. B. (eds), pp. 297–304. Springer.

Derflinger, G., Hörmann, W., and Leydold, J. (2010). Random variate generation by numerical inversion when only the density function is known. *ACM Transactions on Modeling and Computer Simulation* **20**(4), 1–25.

Furrer, R., Nychka, D., and Sain, S. (2010). fields: tools for spatial data. R package, version 6.3.

Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. Springer.

Glasserman, P. (2005). Measuring marginal risk contributions in credit portfolios. *The Journal of Computational Finance* **9**(2), 1–41.

Gross, J. (2006). nortest: Five omnibus tests for the composite hypothesis of normality. R package, version 1.0.

Gyorey, T., Jochim, M., and Norton, S. (2010). McKinsey global survey results: the challenges ahead for supply chains. Survey, McKinsey and Company.

Haksöz, Ç., and Kadam, A. (2009). Supply portfolio risk. *The Journal of Operational Risk* **4**(1), 59–77.

Haksöz, Ç., and Seshadri, S. (2007). Supply chain operations in the presence of a spot market: a review with discussion. *Journal of Operational Research Society* **58**(11), 1412–1429.

Haksöz, Ç., and Şimşek, K. D. (2010). Modeling breach of contract risk through bundled options. *The Journal of Operational Risk* **5**(3), 3–20.

Hörmann, W., and Sak, H. (2010). t-copula generation for control variates. *Mathematics and Computers in Simulation* **81**, 782–790.

Johnson, M. E., and Tenenbein, A. (1981). A bivariate distribution family with specified marginals. *Journal of the American Statistical Association* **76**(373), 198–201.

Karadağ, D. T. (2008). Portfolio risk calculation and stochastic portfolio optimization by a copula based approach. Master's Thesis, Boğaziçi University Istanbul.

Kole, E., Koedijk, K., and Verbeek, M. (2007). Selecting copulas for risk management. *Journal of Banking and Finance* **31**(8), 2405–2423.

Leydold, J., and Hörmann, W. (2008). Runuran: R interface to the UNU.RAN random variate generators, version 0.8. Department of Statistics and Mathematics, Vienna University of Economics and Business.

Malevergne, Y., and Sornette, D. (2006). *Extreme Financial Risks: From Dependence to Risk Management*. Springer.

Matteis, R. (2001). Fitting copulas to data. Diploma Thesis, Institute of Mathematics, University of Zurich.

Prause, K. (1997). Modelling financial data using generalized hyperbolic distributions. Preprint 48, Center for Data Analysis and Modeling, University of Freiburg.

R Development Core Team (2008). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna. URL: www.R-project.org.

Ricci, V. (2005). *Fitting Distributions with* R. R Foundation for Statistical Computing, Vienna.

Sak, H., and Hörmann, W. (2011). Fast simulations in credit risk. *Quantitative Finance*, forthcoming.

Sak, H., Hörmann, W., and Leydold, J. (2010). Efficient risk simulations for linear asset portfolios in the t-copula model. *European Journal of Operational Research* **202**(3), 802–809.

Wagner, S. M., Bode, C., and Koziol, P. (2009). Supplier default dependencies: empirical evidence from the automotive industry. *European Journal of Operational Research* **199**, 150–161.

Yan, J. (2007). Enjoy the joy of copulas: with a package copula. *Journal of Statistical Software* **21**, 1–21.

Yan, J., and Kojadinovic, I. (2010). copula: multivariate dependence with copulas. R package, version 0.9-7.